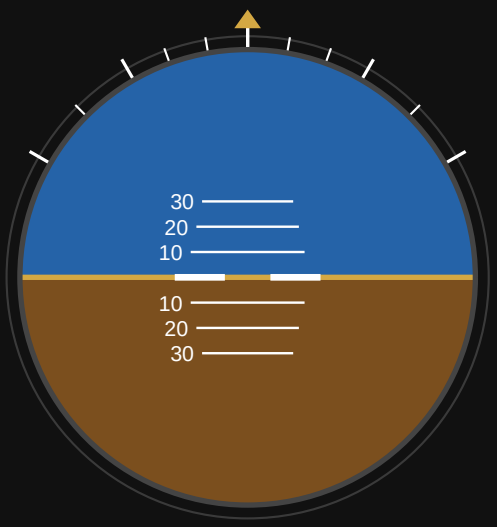

ϕ 0.0° θ 0.0°
Artificial Horizon (ϕ, θ)

ψ 270°
Heading ψ

GPS POSITION DEGRADED

Inertial navigation active · Position error growing

Flying Blind?

What happens when GPS fails –
and what technology does about it

Dipl.-Inform. (FH) Gerd Schmitt
info@sc-sarl.com

Flying Blind?

What happens when GPS fails – and what technology does about it

*For everyone who has ever sat on a plane and wondered:

What actually keeps us on course up here?*

Introduction – A quiet night over the Baltic Sea

It is 3:17 in the morning. Somewhere between Helsinki and Hamburg, at cruising altitude of 11,000 metres, most of the 180 passengers on board are asleep. Those still awake stare out the window into absolute blackness – no lights, no horizon, only the blinking position lights of their own wing reflecting in the glass.

In the cockpit, the calm concentration of a long night sector prevails. The First Officer is drinking his third coffee. The Captain routinely checks the instruments.

Then, without warning, without a bang, without any dramatic event: a small symbol on the navigation screen changes colour. Yellow. A discreet message appears: *GPS POSITION DEGRADED.*

The GPS signal is gone.

Not failed like a light going out. Simply – unreliable. Somewhere down below, on an aircraft carrier or in an inconspicuous vehicle parked on a coastal road, a device is transmitting radio signals that overpower the GPS satellite signals. A jammer. Deliberate. Precise. Effective.

What happens now?

That is the question this book revolves around. Not as a thriller – although the story does have tension. But as an honest, understandable answer to something that affects millions of air travellers and that appears in the news with increasing frequency, without anyone really explaining what lies behind it.

This book explains it. Without formulae in the running text. Without an engineering degree as a prerequisite. But also without false simplifications that obscure more than they illuminate.

By the end you will understand what really happens on that night over the Baltic Sea – technically, humanly, and politically. And you will have an answer to the question that has probably been dormant in you for some time: *How dangerous is this, really?*

The answer is: it depends. And what it depends on is exactly what this book is about.

Chapter 1 – What GPS Really Is (and What It Is Not)

The misunderstanding almost everyone has

Ask ten people on the street what GPS does, and nine will answer: "GPS navigates." It shows where you are. It calculates the route. It tells you when to turn.

All of that is true – but it actually describes the app on your phone, not GPS itself.

GPS – the *Global Positioning System* – does something far more modest and at the same time more elegant: it delivers four numbers. Latitude. Longitude. Altitude. And the precise time. That is it.

What your phone, your navigation system, or an aircraft autopilot makes of that – the route, the map, the course – is the performance of the software. GPS itself is merely a highly precise, worldwide time reference and position provider.

Why does this matter? Because you can only understand what is lost in a GPS failure if you have understood what GPS actually delivers.

How GPS works – in three minutes

Imagine you are standing in the middle of a vast hall and want to know exactly where you are standing. Hanging from the ceiling are thirty loudspeakers, all broadcasting their name simultaneously and giving a precise beat along with it.

If you hear that loudspeaker A's beat reaches you 0.0001 seconds earlier than loudspeaker B's, then you know: you are standing closer to A than to B. With enough loudspeakers and precise clocks, you can calculate your position to within a few centimetres.

That is exactly what GPS satellites do. There are roughly thirty active satellites orbiting the Earth at about 20,000 kilometres altitude. Each one continuously transmits a signal that essentially says: *I am satellite No. 14, and it is now exactly 03:17:42.000000.*

Your GPS receiver listens to at least four of these signals simultaneously. From the tiny time differences with which the signals arrive, it calculates its position. Three satellites would suffice for a position in a plane – the fourth resolves the ambiguity in altitude and corrects the inaccuracy of the receiver's own clock.

The result: position accurate to about three to five metres, worldwide, at any time, in any weather. A masterpiece of engineering that has become so commonplace we take it for granted.

What GPS does not deliver

Here comes the critical limitation that is missing from most explanations:

GPS tells you where you are. It does not tell you how you feel.

That sounds strange for a technical system – but it hits the core. A GPS receiver knows nothing about whether the aircraft is currently rolling, pitching, or climbing. It does not know which direction the nose is pointing. It does not know whether the aircraft is sitting level or being shaken by turbulence.

It only knows the position – and if you compare two successive positions, you can derive the speed and roughly the direction. But that is a derivation, not a direct measurement.

For an autopilot or a pilot who wants to keep the aircraft stable, that is not enough. He needs at every moment the *attitude* of the aircraft in space: how far is the left wing tilted downward? Is the nose rising or sinking? Is the aircraft turning right?

This information GPS simply does not provide. For that, there is an entirely different technology – one that was flying before the first GPS satellite was ever launched.

What GPS failure really means

When GPS fails over the Baltic Sea, the aircraft does not lose its horizon, its controls, its attitude. It loses the precise position information.

That is a difference that counts. An aircraft without GPS is not blind. It is – to stay with the analogy – like an experienced hiker who has forgotten his GPS watch. He still knows how he walks, how he holds himself, how fast he is moving. He just no longer knows quite exactly where on the map he currently is.

How long he can navigate before getting lost depends on what other instruments he has. And that is exactly what the following chapters are about.

Chapter 2 – The Aircraft Feels Itself

Navigation without satellites – an old art

On 21 February 1921, a mail aircraft took off from New York and landed – without any external assistance, without radio, without a ground station – at the planned destination in San Francisco. The pilot navigated with map, compass, stopwatch, and a feel for wind and speed. Dead reckoning, it is called: you know where you started, you know course and speed, and you calculate where you must be.

This method works. But it accumulates errors. Every small inaccuracy in the course, every underestimated gust of wind, every misread speed adds up. Over short distances the errors are negligible. Across the Atlantic they can amount to hundreds of kilometres.

For transcontinental and oceanic aviation, something better was needed. It was found in a physical principle that Isaac Newton already knew in the 17th century, but which could only be realised as a practical navigation system in the 1950s: inertial navigation.

Inertia as a navigation principle

Inertia is the tendency of a body to maintain its state of motion. A rolling billiard ball goes straight until a force deflects it. A flying aircraft maintains its course until engines, rudders, or wind change something.

The interesting part: if you know the acting forces, you can calculate how the motion changes. And forces can be measured – with an accelerometer.

A gyroscope additionally measures the rotation rate: how quickly is the aircraft rotating about its own axis? Is it rolling to the left? Is the nose pitching down?

Together, three accelerometers and three gyroscopes – one each per spatial axis – form an *Inertial Measurement Unit*, abbreviated IMU. With it, a computer can calculate at every moment:

- How is the aircraft currently tilted?
- How quickly is this tilt changing?
- In which direction is it accelerating?

And if you know all of that and also know where you started, you can calculate position and attitude at any time – entirely without GPS, entirely without external signals.

From mechanical gyroscopes to the silicon chip

Early inertial navigation systems were mechanical marvels. Precision gyroscopes, mounted in gimbal suspensions, maintained their orientation in space regardless of how the aircraft moved. These systems were large, heavy, expensive, and sensitive. They cost more than a small car, and a single grain of sand in the wrong bearing could ruin the precision.

Today things look different. Thanks to the miniaturisation of semiconductor technology, gyroscopes and accelerometers exist that sit on a chip no bigger than a fingernail. They are called MEMS sensors – *Micro-Electro-Mechanical Systems*. Tiny mechanical structures, etched in silicon, as thin as a human hair.

How a MEMS gyroscope works

Imagine a tiny tuning fork – two vibrating tines, brought precisely into resonance. When the tuning fork rotates, the Coriolis force deflects the tines out of their plane of vibration. This tiny deflection is measured electrically. The result: a rotation rate, precise to fractions of a degree per second.

The same principle that was once reserved for gyroscopes with diamond bearings now sits in every smartphone – so the screen knows whether you are holding the device landscape or portrait.

A modern MEMS sensor costs a few euros, weighs grams, and delivers measurements at a rate of a thousand times per second. It fits on a circuit board that fits in an envelope.

That sounds like a perfect solution. And it would be – if there were not a fundamental problem that is deeply rooted in physics and that no amount of engineering ingenuity can completely solve.

The unavoidable problem: noise and drift

No sensor measures perfectly. Every measurement contains noise – small, random deviations from the true value. For a thermometer that shows 19.8°C instead of 20.0°C, that is harmless.

For a gyroscope it is a different story.

The strapdown algorithm – the procedure by which a computer calculates attitude and position from sensor measurements – *integrates* the rotation rate. That means: it adds up every single measurement over time. What sounds small has an unpleasant property: errors accumulate.

If the gyroscope makes a random error of 0.01 degrees every second, then after 100 seconds the accumulated attitude error is already 1 degree. After 1,000 seconds it is 10 degrees. The aircraft believes it is flying straight ahead – but the calculated attitude deviates ever further from reality.

Even worse is the *bias* – a systematic error that does not scatter randomly but pulls constantly in one direction. A gyroscope with a tiny bias of 0.1 degrees per hour makes an error of 1 degree in attitude after 10 hours. And this error grows relentlessly without GPS correction.

An analogy

Imagine navigating through a city with a blindfold on. A friend whispers to you every two seconds: "Now ten steps straight ahead. Now turn left, 90 degrees. Now five steps." You follow his instructions exactly.

But your friend makes a small error each time: his "90 degrees" might be 91 degrees. His "ten steps" are nine and a half. Each error in itself is trivial. But after twenty turns you are hundreds of metres away from the planned route – without knowing it.

That is inertial navigation without a correction signal.

Exactly here lies the crux. An IMU alone cannot navigate an aircraft precisely over a long period. It can keep it stable in the short term – very well, in fact. But without an external correction signal that keeps resetting the growing error back to zero, the solution drifts.

GPS delivers exactly this correction signal. It updates the position four times per second and prevents the error from growing.

When GPS fails, the clock starts ticking.

Chapter 3 – The Problem with Time: Why Errors Grow

One minute is harmless. One hour is dangerous.

When GPS fails over the Baltic Sea, little is noticeable in the first minute. The aircraft flies stably. The attitude – roll, pitch, yaw – is still known very precisely, because the inertial sensors handle that excellently in the short term. The position drifts a little, but still in harmless magnitudes.

After five minutes the position error is already noticeably larger – on the navigation display, one to two kilometres of deviation might be showing.

After half an hour we are talking about double-digit kilometre errors if no other correction takes place.

After one hour, inertial navigation is largely unusable for precise purposes.

This time curve is not a weakness of cheaper systems. It is a fundamental physical phenomenon. Even the most expensive military inertial navigation systems, costing six-figure sums, drift eventually. The only question is how quickly.

Why errors do not simply stay small

The reason lies in the mathematics of integration. Whoever still has the word integration in their head from school can use it – for everyone else, here is the essence:

When a gyroscope measures the rotation rate in order to calculate the current angle, it must add up all previous measurements. Imagine wanting to know how far you have walked – not by directly measuring the distance, but by counting and adding up each step. That works. But every counting error accumulates.

Now, to calculate the position, one more step is needed: the acceleration is integrated to obtain the velocity. And then the velocity is integrated once more to obtain the position. Two successive accumulations.

This has an unpleasant property: errors do not grow linearly, they grow with the square of time. A small error in the acceleration measurement, still completely harmless after one minute, can be a hundredfold larger after ten minutes.

The image of the leaking bucket

Imagine a bucket with a tiny hole. It loses one litre per hour. After one hour, one litre is missing – barely noticeable. But the hole does not get smaller; it stays the same. After ten hours, ten litres are missing. The deficit grows steadily.

Now imagine the hole grows over time. That is the behaviour of navigation errors under double integration. The error grows and grows – and increasingly quickly.

What this means in practice: three scenarios

Scenario 1: Short failure, 30 seconds.

The GPS signal drops briefly and returns – perhaps when overflying a heavily overcast area or a brief jamming zone. Inertial navigation bridges this without difficulty. Position errors in the single-digit metre range. Barely relevant for the pilot. Systems like these are designed for exactly such short interruptions.

Scenario 2: Medium failure, 2–5 minutes.

GPS stays gone, perhaps in a jamming area. Inertial navigation keeps the aircraft stable and delivers a usable but increasingly inaccurate position. Deviations in the range of a few hundred metres to a kilometre. The pilot knows he is in a corridor – but where exactly in that corridor becomes increasingly uncertain. Non-critical for cruising over open water. Near mountains or on approach to an airport it begins to get uncomfortable.

Scenario 3: Long failure, 10 minutes and more.

Now it gets serious – not for the aircraft itself, but for the precision of navigation. Errors in the kilometre range. The airport is still safely reachable, but precise approach routes can no longer be maintained. Switching to other navigation sources is now not optional but mandatory.

The good news: GPS is not the only source

Anyone who has read this book to this point might have the impression: without GPS, a modern aircraft is lost.

That is not true. It is not true in several respects.

Firstly, modern commercial aircraft have redundant navigation systems. Multiple independent IMUs, whose errors can be statistically weighed against each other. High-precision inertial

navigation systems that drift considerably more slowly than cheap MEMS chips. And a whole range of other aids.

Secondly, there are ground-based navigation aids – VOR beacons, DME distance measurers, ILS landing aids – that work entirely independently of GPS and are distributed across the whole world. While some of these systems are thinly spread in remote areas, on the main traffic routes across Europe and the North Atlantic the coverage is excellent.

Thirdly – and this is the actual subject of this book – there are modern filter algorithms that extract the best from all available sources and radically reduce the error while GPS is down.

How this works is explained in the next chapter. And it is, honestly, one of the most beautiful examples of how mathematics can elegantly solve a practical problem – without needing to know a single symbol to understand it.

End Part A (Introduction, Chapters 1–3)

Chapter 4 – The Sceptical Navigator

Two witnesses, two statements, one truth

Imagine the following situation: you are a detective investigating exactly where a suspect was at midnight. You have two witnesses.

Witness A is precise but unreliable. He says: "He was standing exactly on the second paving stone from the left, third row." But this witness has a known weakness: he sometimes misremembers, and the longer he thinks about it, the further his account drifts from the truth.

Witness B is imprecise but consistent. He says: "He was standing somewhere near the lamp post, roughly here." He points to an area of about ten square metres. Imprecise – but constant. Witness B knows what he knows and does not change his account.

Which do you believe? Trusting either one alone would be a mistake. The smart answer is: you combine both accounts. The precision of Witness A, weighted with the knowledge that he becomes less reliable over time. The consistency of Witness B, weighted with his inherent imprecision.

That is exactly what the Kalman filter does. It is the sceptical but fair detective of navigation.

What the Kalman filter really is

Rudolf Kálmán was a Hungarian-American mathematician who in 1960 published a paper that changed applied mathematics forever. His filter – an algorithm that optimally combines multiple uncertain measurement sources – became so quickly so important that it enabled the Apollo Moon landing programme. Without the Kalman filter, Neil Armstrong would not have landed on the Moon in 1969.

Today the Kalman filter is in every smartphone, every navigation system, every modern aircraft. It is one of the most widely implemented algorithms in history.

Its core idea can be summarised in one sentence:

Trust no source blindly. Weight each source according to its current reliability – and update this weighting with every new piece of information.

The sceptical navigator – an analogy

Imagine an experienced navigation officer on a ship sailing through dense fog. He has two aids:

First, his logbook. It records where the ship was an hour ago, how fast it was travelling, and in which direction. From this data he can calculate where the ship *should* be now. This is his prediction – it is good, but not perfect, because currents and small course corrections are not fully captured.

Second, a radar that occasionally detects a coastal profile through the fog. This measurement is blurred – the fog distorts the image – but it is independent information that has nothing to do with the logbook.

What does the navigator do? He compares his prediction with the radar measurement. If both agree: good, he trusts his prediction. If they diverge: he weights them. How reliable is the radar right now? How old is the last logbook entry? He shifts his estimate towards the more trustworthy source.

And then – this is the decisive step – he updates his logbook with the new combined estimate. At the next comparison he does not start from scratch. He learns from every measurement.

That is the Kalman filter. Nothing more and nothing less.

In the aircraft: IMU and GPS as two witnesses

Applied to our aircraft:

The IMU is Witness A: highly precise, fast, delivers a measurement a thousand times per second – but accumulates errors over time that grow and grow.

GPS is Witness B: delivers a measurement only four times per second, is slower, can be jammed – but does not make errors that accumulate. Every GPS measurement is independent of the previous one.

The Kalman filter combines both. At every moment it knows:

- How long has the IMU been running without correction? (If long: trust it less)
- How good is the GPS signal right now? (If poor: trust GPS less)
- What is the most probable truth when both are combined with their respective weights?

The result is better than either source alone.

What happens when GPS fails?

When GPS fails, the Kalman filter loses one of its witnesses. It continues working – but now only with the IMU. And it *knows* that it only has one witness. It knows that its estimate is becoming less certain.

Modern implemented systems carry internally a number one can think of as a "confidence radius": within what range does the true position probably lie? This radius grows when GPS fails – slowly at first, then faster.

The pilot sees on his display not just a position, but also an indication of how reliable that position is. A conscientious system shows him: "I know my position – but only to within 500 metres now." And five minutes later: "Only to within two kilometres now."

That is not a failure of the system. That is the system honestly communicating what it knows and what it does not. And honest uncertainty statements are worth their weight in gold in navigation.

How good can it get?

Here come the numbers from a real project that forms the basis of this book: in a flight simulator, a modern GPS/IMU navigation system with a Kalman filter was built and tested. The system was subjected to a controlled GPS failure lasting 150 seconds – two and a half minutes.

The result:

- **Without Kalman filter:** 1,258 metres position error after 150 seconds.
- **With Kalman filter:** 37 metres position error after 150 seconds.

Factor 34. Not as a theoretical value – as a measured result, reproducible, confirmed in twenty different test scenarios.

37 metres. That is less than the length of a football pitch. After two and a half minutes of GPS failure.

That is the difference a well-implemented Kalman filter makes.

Chapter 5 – What Happens in the First 30 Seconds?

Back to the night over the Baltic Sea

The GPS symbol lights up yellow. *GPS POSITION DEGRADED.*

At this moment, several things happen simultaneously – in the systems, in the cockpit, and in the minds of the pilots.

The first three seconds: automatic switchover

Modern commercial aircraft do not have one GPS but several. An Airbus A320, for example, operates three independent inertial reference systems (IRS) and can use multiple GPS receivers. These systems work continuously in parallel and compare their results.

If one of the GPS signals drops out or becomes implausible, the systems switch automatically within the first few seconds. The *Flight Management System* (FMS) – the onboard computer that plans and monitors the flight – registers the failure and immediately begins estimating navigation accuracy from several remaining sources.

The pilot sees a warning indication. But the aircraft keeps flying. Nothing wobbles. Nothing tilts. The autopilot holds the course.

Seconds 10 to 30: establishing situational awareness

Now the pilot is called upon – not as a response to an emergency, but as the routine execution of trained procedures. Checklists exist for exactly this situation. They are trained, regularly practised, and in practice so internalised that an experienced captain can run through them half asleep.

What happens in concrete terms? The pilot establishes which navigation systems are still working. He checks whether the inertial reference systems (IRS) are unaffected – they are, because GPS jamming does not affect the IRS. He checks whether ground-based navigation sources are available: VOR beacons, DME distance measurers.

Over the Baltic Sea, on a main traffic route, these signals are almost always present. Navigation is restricted, but far from blind.

What the autopilot does – and what it no longer does

The autopilot of a modern commercial aircraft is not a single system but a network of several control modules that interlock. Simplified, there are two levels:

The **attitude control** keeps the aircraft stable: roll, pitch, yaw are controlled by the IMU data. This works completely independently of GPS and continues as if nothing had happened.

The **course-keeping and navigation** – the ability to follow a planned flight path – depends on position information. Here the autopilot notices the difference. It can still hold a course (fly a direction), but it can no longer follow a planned waypoint to the metre when the position becomes uncertain.

In practice this means: the autopilot keeps the aircraft airborne and on roughly the right course. But the fine precision is lost.

The human side: training and composure

What perhaps surprises most in this scenario: for well-trained pilots, GPS failure is not an emergency. It is an unusual situation – but one for which years of training exists.

Commercial pilots are regularly confronted with system failures in simulators. GPS failure is one of the more benign scenarios. Engine failure, decompression, fire on board – those are the scenarios that draw sweat.

GPS failure over open water, on a known route, with functioning inertial reference systems? That is a routine malfunction. Uncomfortable, but manageable.

The composure in the cockpit in such moments is not accidental. It is the result of hundreds of hours in the simulator, in which exactly such situations were practised until the responses run automatically.

What passengers feel

Nothing. Absolutely nothing.

Perhaps a slight pause before an announcement comes – or no announcement at all. GPS failure is in the vast majority of cases not an event that affects passengers. The aircraft keeps flying. The drinks keep being served. The peanuts keep landing in the little plastic bowl. The flight lands with high probability on time at the destination.

Unless – and this is the point that Chapter 6 will address – the jamming is not short and random, but deliberate, persistent, and strategically positioned.

When does it get really dangerous?

There are situations in which GPS failure is considerably more serious than the cruise-flight scenario described above.

The approach to an airport is one of them. Modern precision approaches – particularly in poor visibility – rely on GPS-based navigation systems (GNSS approaches). If GPS fails on approach, a switch must be made to less precise procedures. That means: higher weather minima, possibly diverting to another airport.

Helicopters in mountain areas or on offshore platforms operate in environments where errors of just a few metres can make the difference between a safe approach and a catastrophe. GPS failure is here considerably more critical than at cruising altitude over the Baltic Sea.

Unmanned aerial vehicles (drones) are often entirely dependent on GPS, because they have no human perception as a backup. For them, GPS failure is acutely dangerous.

But for an airliner at cruising altitude? The aircraft stays in the air. It remains navigable. And if the failure lasts longer, there are protocols.

Chapter 6 – The Baltic Sea Is No Coincidence

GPS as a military system

GPS is not a civilian system. It was developed in the 1970s by the US Department of Defense and remains under military control to this day. The US could shut down or degrade the system for civilian users at any time – although they committed in 2000 to provide it free of charge permanently.

This historical fact has an important consequence: whoever jams GPS is not merely attacking civil aviation. They are attacking a military navigation system.

And whoever knows that understands why GPS jamming in the Baltic is not a harmless phenomenon.

Jamming and spoofing – two very different threats

There are two fundamentally different ways to disrupt GPS, and it is worth knowing the difference.

GPS jamming is broadband overloading: a jammer floods the GPS frequency with noise. The GPS device can no longer filter out the real satellite signals and loses the position. It knows it has no signal – the display shows "no GPS", and the pilot switches to other navigation means.

Jamming is blunt but effective. Within a radius of a few dozen to a few hundred kilometres around a medium-strength jammer, it can knock out all GPS receivers. And a jammer does not need to be a high-tech device: for years, instructions for jammers costing under 100 euros to build circulated in online forums.

GPS spoofing is considerably more sophisticated – and considerably more dangerous. A spoofer does not transmit noise but forged GPS signals. Signals that look like genuine satellite signals but contain false position data.

The GPS device notices nothing. It receives signals that look technically correct and calculates a position – the wrong one. The aircraft believes it is flying over the sea while it is in fact heading towards a mountain. Or it believes it is at 10,000 metres altitude while the calculated position is already at ground level.

Spoofing attacks on civilian aircraft have been documented over the eastern Mediterranean and over the Baltic. In some cases aircraft were "shifted" to positions in the middle of military

airfields – presumably as a test or to cause confusion.

The Baltic situation

Since 2022 – in parallel with the escalation of the war in Ukraine – GPS disruptions over the Baltic have increased considerably. The cases are documented, recorded by aviation safety authorities, and intensively discussed in professional circles.

The disruptions concentrate on areas near the Russian exclave of Kaliningrad and along the Russian-Finnish border. The range is considerable: disruptions have been reported at high cruising altitude over parts of Finland, Estonia, and Latvia.

What lies behind it? Military GPS jamming systems that Russia uses for the protection of its own systems radiate in a range that extends far into civil airspace. Whether this is deliberate or a side effect of military operations remains diplomatically contested. Technically the effect is the same.

The response of aviation

The international aviation organisation ICAO and the European Aviation Safety Agency EASA have responded to the increase in disruptions – albeit at the usual bureaucratic speed. Pilots are informed in briefings of possible GPS impairments. Airlines have updated internal procedures.

What has not happened: a fundamental technical solution. Aviation has become GPS-dependent, and this dependency cannot be reversed overnight. Ground-based navigation systems such as VOR have been decommissioned in many places in recent years for cost reasons – at exactly the time when GPS disruptions are increasing. A suboptimal development.

The radical solution

There is of course also the direct approach. A fighter jet accompanies the passenger aircraft, and when a GPS disruption occurs, a nuclear-armed cruise missile is immediately fired at the position of the jammer. After the detonation, GPS should be working again.

Technically impeccable. Success rate: approaching 100 percent. Cost-benefit ratio from a purely navigational standpoint: remarkably favourable.

The ICAO does not recommend this procedure for reasons as yet unexplained. But it illustrates a genuine point: in a military conflict, GPS jamming would provoke a completely different category of response than in civil aviation. Military systems possess means of neutralising jamming sources that civilian passengers are better off not thinking about while eating their peanuts.

> In the airliner over the Baltic, the answer remains more modest – and more sensible: better filters, redundant systems, and well-trained pilots.

The real vulnerability

The most honest assessment of GPS vulnerability reads: the problem lies not in GPS technology itself, but in monoculture.

Modern aviation relies too heavily on a single technology. GPS is good, GPS is precise, GPS is cheap. And for exactly that reason aviation – like agriculture with monocultures – has created a highly efficient but vulnerable infrastructure.

The solution is not a single new technology but redundancy: multiple independent systems that check and complement each other. GPS plus inertial navigation plus ground-based radio navigation plus – little used in the civilian sector but available – celestial navigation and barometry.

No single system is safe. A well-combined composite system is very robust.

End Part B (Chapters 4–6)

Chapter 7 – What Technology Can Do Today – An Honest Assessment

Not a marketing brochure

This chapter is not advertising. It is a sober performance report – based on a real navigation system that was built from scratch, implemented, and validated in a flight simulator with twenty regression tests as part of a development project.

The numbers are real. The limits are real. And the honesty about what does not work is – in the realm of technology communication – unfortunately rarer than it should be.

The test system: what was built?

The system underlying this assessment is a navigation system for a light aircraft, built from off-the-shelf MEMS sensors. The hardware:

- **Gyroscope and accelerometer:** ICM-42688, a modern MEMS chip for about 15 euros. Measures rotation rate and acceleration on three axes, 1,000 times per second.
- **Magnetometer:** MMC5983MA – measures the Earth's magnetic field for heading determination.
- **Barometric altitude sensor:** BMP390 – measures air pressure, from which altitude is derived.
- **GPS receiver:** Standard module, four position updates per second.

Total hardware cost: under 50 euros. That is the point that would have been unimaginable twenty years ago.

The software: a self-developed Kalman filter with 16 state variables – meaning 16 quantities estimated simultaneously and continuously corrected. These include position, velocity, attitude, and also the errors of the sensors themselves (bias, drift) – the system thus estimates not only where the aircraft is, but also how wrongly its own sensors are currently measuring.

The decisive test: 150 seconds of GPS failure

The most informative test was the GPS failure test. The aircraft flies a traffic pattern – a standardised rectangular circuit around an airfield, roughly 3 kilometres × 1.5 kilometres. During the flight, GPS fails for exactly 150 seconds.

After these 150 seconds: how far is the estimated position from the true position?

Result without Kalman filter (raw inertial navigation only):

1,258 metres position error.

That corresponds to roughly twelve football pitches. The aircraft would believe it was in a completely different location.

Result with Kalman filter:

37 metres position error.

Less than the length of a football pitch. After two and a half minutes of GPS failure.

That is an improvement factor of 34. Not as an ideal value under perfect conditions – as a robust result confirmed in twenty different test scenarios.

What makes the difference

The Kalman filter achieves so much more because it does several things simultaneously:

It estimates the sensor errors as well. That sounds trivial but is decisive. A gyroscope has a bias – a systematic error pulling in a certain direction. If the filter knows this bias and accounts for it, it can remove it. While GPS is still running, the filter "learns" the bias of the sensors. When GPS fails, it applies this knowledge.

It uses all available measurements. Barometer, magnetometer, airspeed sensor – all deliver information that the filter incorporates. No source is perfect, but together they produce a more consistent picture.

It knows how uncertain it is. This self-awareness is perhaps the most important property. A system that knows its own uncertainty can communicate honestly how trustworthy its output is. That is critical for safety.

The limits – clearly stated

Now the uncomfortable truths.

150 seconds is not 30 minutes. The test ran for two and a half minutes. With longer GPS failure, the error continues to grow – more slowly than without a filter, but it grows. After ten minutes, the error would be considerably above 37 metres. After one hour we would be in ranges that are unusable for precise navigation.

Cheap MEMS sensors have limits. The described setup for under 50 euros achieves something remarkable. But a professional inertial navigation system costing 50,000 euros drifts considerably more slowly. The quality of the hardware sets the limit of what software can correct.

Magnetometers are susceptible to interference. Metal nearby, electrical cables, interference fields – all of these corrupt the heading determination by the magnetometer. In an aircraft full of electrical systems, that is a real challenge.

The filter is only as good as its model. A Kalman filter is based on a mathematical model of the system. If the aircraft does something not foreseen in the model – for example extreme manoeuvres, stall, turbulence – the filter can become unreliable.

What this means in practice

For an airliner at cruising altitude, the described performance range is more than sufficient. A position error of 37 metres after two and a half minutes of GPS failure is irrelevant for route navigation at 10,000 metres. The aircraft is still very far from the nearest obstacle.

It becomes critical only when:

- the GPS failure lasts longer than five to ten minutes,
- you are near mountains or other obstacles,
- a precision approach must be flown.

In all these cases there are fallback procedures. No modern commercial aircraft is built so that it becomes headless when GPS fails.

Chapter 8 – What the Pilot Can Do, What He Cannot – and What Protects Him

The human in the system

In recent decades aviation has automated itself with impressive consistency. Modern aircraft can take off, hold course, descend, and land – all fully automatically. In some operating modes the pilot sits there primarily to monitor the system and intervene when something unexpected happens.

GPS failure is exactly such a moment: a moment in which the automation is limited and the pilot steps more to the forefront again. And it is a moment for which pilots are explicitly trained.

What a well-trained pilot does in this situation

The first reaction is not panic but situational awareness. The pilot establishes what he has and what he does not have.

He has: functioning engines, stable attitude, functioning inertial navigation, ground-based radio navigation (VOR/DME), radio contact with air traffic control.

He does not have: precise GPS position.

From this inventory, clearly defined steps follow: inform air traffic control. Switch to alternative navigation. Hold altitude. Hold course.

That sounds simple – and in a normal situation it is. The art lies in executing these simple steps correctly even when other things are happening simultaneously: bad weather, tired crew, nervous passengers.

What protects the pilot: the system behind the system

Individual pilots make mistakes. That is not a criticism – it is a fact about humans. That is why modern aviation is constructed so that a single mistake rarely leads to an accident. It typically takes a chain of errors.

The tools that build this safety chain:

Two pilots in the cockpit: Not because one could not fly the aircraft alone, but because two eyes and two brains catch errors that one alone would make. One flies, one monitors. In case of doubt: cross-check. In case of discrepancies: discussion.

Checklists: In aviation, checklists are not tedious bureaucracy but a cognitive safety net. They ensure that no step is forgotten even under stress. The checklist for GPS failure is short, clear, and practised.

Air traffic control: ATC knows the traffic picture. On a GPS failure report it can use secondary radar (radar that is not GPS-dependent) and give the pilot radar vectors – direct course instructions based on ground radar. In densely occupied airspace like Europe, that is a reliable backup.

Minimum safe altitudes (MSA): No aircraft may fly below a defined altitude in a defined area – even if the pilot wanted to. These minimum altitudes take into account all obstacles in the vicinity. As long as the aircraft remains above the MSA, there is no collision risk with the ground.

The honest answer to the question: how dangerous is it?

For an airliner, at cruising altitude, over the Baltic, with an experienced crew and modern aircraft:

GPS failure is not an emergency.

It is an unusual situation that demands attention and in which protocols are executed. But the probability that this failure leads to an accident is vanishingly small. The redundant systems, the trained pilots, and air traffic control form a safety net that is extraordinarily more robust than newspaper articles about GPS disruptions suggest.

That does not mean GPS disruptions are harmless. They are a serious problem:

- They increase pilot workload.
- They can lead to irregularities in arrivals and departures.
- In extreme cases they can force diversion decisions that are costly and disruptive.
- For smaller, less well-equipped aircraft the risk is greater.

But a dark night over the Baltic, GPS gone, pilots asleep – this image is fiction. In reality the pilots are not asleep, the systems keep working, and the training kicks in exactly when it is needed.

What needs to change – and what is already changing

The aviation community has recognised GPS vulnerability as a structural problem. The answers are still emerging, but they are real:

More redundancy: New aircraft are equipped with multiple independent navigation systems that check each other.

Better filters: Next-generation Kalman filters estimate not only the sensor biases but model the entire error profile of the system over time. Improved algorithms can compensate for longer GPS outages.

Ground-based infrastructure: Several European countries are discussing the reactivation of eLoran – a long-wave navigation system that is independent of GPS and practically unjammable.

International diplomacy: GPS jamming in the Baltic is now regularly on the agenda in ICAO committees and bilateral discussions. Short-term solutions are not in sight – but the problem is named.

Epilogue – Back over the Baltic Sea

The scene gets its ending

It is 3:51. The GPS symbol remains yellow. In the past thirty-four minutes, the crew has routinely switched to ground-based navigation, informed air traffic control, and held the course with slightly heightened attention.

The First Officer has finished his coffee. The Captain has spoken with Helsinki ATC twice.

At 3:58, the symbol switches back to green. *GPS POSITION ACQUIRED.* The FMS immediately accepts the updated position. The deviation against the internally calculated inertial navigation: 1.2 kilometres. For a flight at 11,000 metres cruising altitude: irrelevant.

At 5:23, the flight lands in Hamburg. On schedule.

What the passengers experienced: a smooth flight. Some turbulence over the Danish coast. And in row 27, a slightly too-loud snorer.

What the crew experienced: a GPS disruption that is in the manual, that was trained for, and that passed without incident.

What is really at stake

The story of this night ended well – and it almost always ends well. But that is no reason for complacency.

GPS disruptions are increasing. The dependency of civilian infrastructure on GPS is also increasing – not just in aviation, but in logistics, finance (GPS timestamps matter for transactions), telecommunications, and agriculture.

A society that builds a critical infrastructure on a single, easily jammable system creates a vulnerability. Not for the normal case – the normal case works. But for the disruption case that comes when someone wants to exploit this vulnerability.

The technology described in this book – Kalman filters, sensor fusion, redundant navigation – is part of the answer. But only part. The other part is political: the willingness to invest in redundancy, even when it is more expensive, even when GPS is currently working.

One last number

37 metres.

That is the position error that a well-implemented GPS/IMU system with Kalman filter makes after two and a half minutes of GPS failure. With hardware that costs together less than a good dinner.

37 metres. In a world that navigated across kilometres before GPS was invented.

That is what technology can achieve when it is carefully built, honestly tested, and soberly assessed.

The aircraft over the Baltic still has 37 metres of margin. It knows where it is. It knows how uncertain it is. And it keeps flying.

Appendix – For the Curious

This appendix is optional. It is aimed at readers who wish to explore further. The book is complete without it.

A – How a Kalman filter really calculates

At its core, the Kalman filter operates in two alternating steps that repeat continuously:

Step 1: Prediction

The filter takes its current state of knowledge – where is the aircraft, how is it moving – and calculates where it will be a tenth of a second later. This prediction is not perfect: it is based on a model of the aircraft, and models are always simplified. So the filter also remembers how uncertain its prediction is.

Step 2: Update

A new measurement arrives – for example a GPS position update. The filter compares this measurement with its prediction. If they agree: good. If not: it shifts its estimate a little in the direction of the measurement. How far it shifts depends on how trustworthy it judges the measurement to be (the so-called Kalman gain).

After the update, the filter has an improved estimate. With this it goes into the next prediction. And so on – a thousand times per second.

B – What MEMS means and why it is revolutionary

MEMS stands for Micro-Electro-Mechanical Systems. These are tiny mechanical structures etched in silicon – using the same processes that also produce computer chips.

A MEMS gyroscope contains a microscopic vibrating structure, smaller than a human hair. When the circuit board rotates, the Coriolis force acts on this structure. This force is measured electrically as a change in capacitance – and from this the rotation rate is calculated.

The revolution lies in scaling: what was once an expensive, sensitive precision device now costs a few euros and sits in every smartphone. This has not only democratised navigation, but also enabled robotics, toy drones, fitness trackers, and dozens of other applications.

C – Why quaternions are better than Euler angles

To describe the attitude of an aircraft, three numbers are needed: roll (lateral tilt), pitch (nose up/down), and yaw (direction change). These three numbers are called Euler angles.

That sounds simple – and for display in the cockpit it is. But for calculation in the computer this representation has a fatal weakness: gimbal lock.

Gimbal lock occurs when the aircraft points exactly vertically upward or downward. In this position, roll and yaw become indistinguishable – the mathematical system collapses. For normal aircraft this is rarely relevant, but for fighter jets or model aircraft it is quite possible.

Quaternions are an alternative mathematical representation of attitude that has no gimbal lock. They consist of four numbers instead of three and are more efficient for the computer to process. The disadvantage: they are difficult to understand intuitively for humans. (What does a quaternion of [0.707; 0; 0.707; 0] mean visually? Almost nobody can say that immediately.)

The practical solution: internally the computer calculates with quaternions (precise, efficient, no gimbal lock). For display and for the pilot they are converted to Euler angles (intuitive, understandable). The best of both worlds.

D – Further Reading

For those who wish to explore the subject more deeply:

Popular science:

- "The Signal and the Noise" by Nate Silver – not a navigation book, but an excellent example of how to explain uncertainty and prediction in an accessible way.

Introductory technical:

- "Understanding GPS/GNSS: Principles and Applications" by Elliott Kaplan and Christopher Hegarty – the standard introduction to GPS technology, accessible for technically interested readers without specialist knowledge.
- "Aided Navigation: GPS with High Rate Sensors" by Jay Farrell – solid introduction to sensor fusion and Kalman filtering for navigation, somewhat more mathematical.

For the seriously interested:

• "Optimal State Estimation" by Dan Simon – the standard work on the Kalman filter, mathematically rigorous, with many practical examples.

• "Strapdown Inertial Navigation Technology" by David Titterton and John Weston – the bible of inertial navigation, not for the faint-hearted.

Online:

• The Wikipedia articles on Kalman filter and Inertial Navigation System are unusually good and serve as a first deepening.

E – Glossary of Key Terms

Accelerometer: Sensor that measures acceleration on one or more axes. In MEMS IMUs typically three axes.

Attitude: The orientation of an aircraft in space, described by roll, pitch, and yaw angles.

Bias: Systematic measurement error that pulls constantly in one direction. A gyroscope with positive bias "believes" it is rotating slightly to the right, even when stationary.

Dead reckoning: Navigation by projecting a known starting position forward based on course and speed, without external reference.

DME (Distance Measuring Equipment): Ground-based radio navigation system that measures the slant range to a ground station. Independent of GPS.

ES-EKF (Error-State Extended Kalman Filter): Special Kalman filter variant that estimates not the absolute state but the error in the state. Particularly suited for inertial navigation.

Gimbal lock: Singularity in Euler angle representation that occurs when the pitch angle is exactly ±90°. Solved by using quaternions.

Gyroscope: Sensor that measures rotation rate about an axis. MEMS gyroscopes use the Coriolis effect of a vibrating mass.

ILS (Instrument Landing System): Ground-based landing aid for precision approaches in poor visibility. Independent of GPS.

IMU (Inertial Measurement Unit): Combination of three accelerometers and three gyroscopes on three axes. Measures acceleration and rotation rate of the carrier vehicle.

IRS (Inertial Reference System): Term used in commercial aviation for a high-quality IMU with embedded navigation computer.

Jamming: Deliberate disruption of a radio signal by overpowering it with noise. In GPS jamming, satellite signals are masked by a stronger interference signal.

Kalman filter: Mathematical algorithm for optimal estimation of a system state from multiple uncertain measurement sources. Developed in 1960 by Rudolf Kálmán.

MEMS (Micro-Electro-Mechanical Systems): Miniaturised mechanical structures in semiconductor technology. MEMS sensors are the basis of modern, low-cost IMUs.

NED (North-East-Down): Reference coordinate system widely used in aviation with x-axis pointing North, y-axis East, z-axis Down.

Quaternion: Mathematical representation of a rotation in space by four numbers. No singularities (unlike Euler angles), but less intuitive.

Spoofing: Forged GPS signals that closely resemble genuine satellite signals but contain false position data. More dangerous than jamming because the receiving device does not detect the attack.

Strapdown algorithm: Procedure for calculating attitude, velocity, and position from IMU measurements by numerical integration.

VOR (VHF Omnidirectional Range): Ground-based radio navigation system that provides the magnetic bearing to the ground station. Independent of GPS, in operation for decades.

End of book

Flying Blind? What happens when GPS fails – and what technology does about it

First Edition

Dipl.-Inform. (FH) Gerd Schmitt

info@sc-sarl.com

sc-sarl.com/en/index.html